Efficient multivariate kernels for sequence classification

Pavel P. Kuksa
pavel@pkuksa.org

arXiv:1409.8211v2 [cs.LG] 30 Sep 2014

ABSTRACT

Kernel-based approaches for sequence classification have been successfully applied to a variety of domains, including the text categorization, image classification, speech analysis, biological sequence analysis, time series and music classification, where they show some of the most accurate results.

Typical kernel functions for sequences in these domains (e.g., bag-of-words, mismatch, or subsequence kernels) are restricted to *discrete univariate* (i.e. one-dimensional) string data, such as sequences of words in the text analysis, codeword sequences in the image analysis, or nucleotide or amino acid sequences in the DNA and protein sequence analysis. However, original sequence data are often of real-valued multivariate nature, i.e. are not univariate and discrete as required by typical k-mer based sequence kernel functions.

In this work, we consider the problem of the *multivariate* sequence classification (e.g., classification of multivariate music sequences, or multidimensional protein sequence representations). To this end, we extend *univariate* kernel functions typically used in sequence domains and propose efficient *multivariate* similarity kernel method (MVDFQ-SK) based on (1) a direct feature quantization (DFQ) of each sequence dimension in the original *real-valued* multivariate sequences and (2) applying novel multivariate discrete kernel measures on these multivariate discrete DFQ sequence representations to more accurately capture similarity relationships among sequences and improve classification performance.

Experiments using the proposed MVDFQ-SK kernel method show excellent classification performance on three challenging music classification tasks as well as protein sequence classification with significant 25-40% improvements over univariate kernel methods and existing state-of-the-art sequence classification methods.
Keywords: multivariate sequence classification, string kernels, vector quantization, direct feature quantization, music classification, protein classification

1. INTRODUCTION

Large-scale sequence analysis has become an important task in data mining inspired by numerous applications such as the document and text classification or the analysis of time series, music data, or biological sequences. Classification of string data, i.e. univariate sequences of *discrete* symbols (such as words, amino acids, codewords), has attracted particular attention and has led to a number of new algorithms [4, 11, 19, 14, 29, 23, 33, 34, 3, 37].

In particular, kernel-based approaches for sequence classification show some of the most accurate results for a variety of problems, such as text categorization [14, 38], image classification [23], speech analysis [3], biological sequence analysis [29, 14, 12], time series and music classification [33, 14].

In a kernel-based framework, the classification of the sequence $X = x_1, \ldots, x_{n_X}$ is based on a kernel function $K(X, Y)$ which is computed to measure the similarity between pairs of sequences X and Y. For instance, given a set of positive training instances $\mathcal{C}^+$ and a set of negative training instances $\mathcal{C}^-$, the SVM [32] learns a classification function of the following form

$$f(X) = \sum_{Y^+ \in \mathcal{C}^+} \alpha^+_{Y^+} K(X, Y^+) - \sum_{Y^- \in \mathcal{C}^-} \alpha^-_{Y^-} K(X, Y^-) \quad (1)$$

Typical k-mer-based kernel functions (e.g., mismatch or spectrum kernels, gapped and wildcard kernel functions [17, 22]) essentially rely on *symbolic Hamming-distance* based matching of one-dimensional (1D) k-mers (k-long substrings) in the input sequences. For example, given one-dimensional (1D) sequences X and Y over alphabet Σ (e.g., amino acid sequences with $|\Sigma|$=20), the *spectrum-k* kernel [18] and the *mismatch-*(k,m) kernel [19] measure similarity between sequences as

$$\begin{aligned} K(X, Y|k, m) &= \langle \Phi_{k,m}(X), \Phi_{k,m}(Y) \rangle \\ &= \sum_{\gamma \in \Sigma^k} \Phi_{k,m}(\gamma|X) \Phi_{k,m}(\gamma|Y) \\ &= \sum_{\gamma \in \Sigma^k} \left(\sum_{\alpha \in X} I_m(\alpha, \gamma) \right) \left(\sum_{\beta \in Y} I_m(\beta, \gamma) \right) \\ &= \sum_{\alpha \in X} \sum_{\beta \in Y} \sum_{\gamma \in \Sigma^k} I_m(\alpha, \gamma) I_m(\beta, \gamma) \end{aligned} \quad (2)$$

where

$$\Phi_{k,m}(\gamma|X) = \left(\sum_{\alpha \in X} I_m(\alpha, \gamma) \right) \quad (3)$$

is the number of occurrences (possibly with up to m mismatches) of the k-mer γ in X, and the matching/indicator function $I_m(\alpha, \gamma) = 1$ if α is in the mutational neighborhood $N_{k,m}(\gamma)$ of γ, i.e. α and γ are at the Hamming distance of at most m. This kernel (Eq. 2) essentially amounts to a cumulative Hamming-distance based pairwise comparison of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

all k-mers α and β contained in sequences X and Y, respectively, with maximum number of mismatches m. The level of similarity of each pair of substrings (α, β) here is indicated by the number of identical substrings in the mutational neighborhoods $N_{k,m}(\alpha)$ and $N_{k,m}(\beta)$ of α and β, $\sum_{\gamma \in \Sigma^k} I_m(\alpha,\gamma) I_m(\beta,\gamma)$ (Eq. 2). For the spectrum kernel (m=0), this similarity level is simply the exact matching of α and β.

On the other hand, in many practical applications input sequences are *multivariate*, i.e. input data is in the form of sequences of *R-dimensional* real-valued feature vectors, as opposed to one-dimensional discrete strings. This is the case, for instance, in commonly used MFCC representations for music data as series of 13-dimensional MFCC feature vectors (e.g., [31, 20, 16]) extracted from short time segments, or 20-dimensional profile representations of protein sequences as series of probabilistic amino acid substitution vectors in biological sequence analysis [11, 8].

Such original, multivariate real-valued feature sequences are typically transformed into *univariate sequences* in order to apply a univariate string kernel method such as spectrum or mismatch, e.g., [23, 37, 14, 12]. For example, this transformation is frequently accomplished by applying a vector quantization (VQ) algorithm to feature vectors thus transforming a multivariate R-dim real-valued sequence into a discrete univariate *codeword* sequence.

In contrast, in this work we consider an alternative approach to real-valued multivariate sequence classification which directly exploits these richer multivariate (R-dimensional) sequences (e.g., MFCC feature sequences for music, or sequences of physico-chemical amino acid descriptors for proteins). In this approach, the R-dimensional multivariate sequences are considered as $R \times |X|$ feature - spatio/temporal matrices with rows corresponding to feature dimensions and columns corresponding to temporal or spatial coordinates. Using these representations, we propose an efficient discrete multivariate kernel method (MVDFQ-SK) based on (1) a direct feature quantization (DFQ) (Sec. 3.1) of the original multivariate sequence, and (2) novel manifold-based discrete multivariate kernel functions applied to these discrete DFQ representations (Sec. 3.2, 3.3). The developed approach is applicable to a wide range of sequence domains, both *discrete*- and *real*- valued, such as *music*, *images*, or *biological sequences*.

Experiments using the new multivariate direct feature quantization kernels (MVDFQ-SK) kernels on music genre and artist recognition, as well as protein sequence classification tasks show excellent predictive performance (Sec. 4) with significant *25%-40%* improvements in predictive accuracy over univariate kernel functions and a number of other state-of-the-art sequence classification methods.

2. RELATED WORK

Recently, a large variety of methods have been proposed to solve the sequence classification problem, including *generative*, such as HMMs, or *discriminative* approaches. Among the discriminative approaches, string kernel-based [32, 27] methods provide some of the most accurate results [11, 19, 29, 17, 14, 5, 23, 38] in many sequence analysis tasks.

In the kernel-based approaches, the similarity between sequences X and Y is frequently computed based on the co-occurrence of string features (e.g., k-mers), as in spectrum kernels [18] or substring kernels [34]. Inexact comparison of the sequences in this framework is typically achieved using different families of mismatch [19] or profile [11] kernels. Both spectrum-k and mismatch-(k,m) kernels directly extract string features (k-mers) from the observed sequence, X. On the other hand, the profile kernel, proposed by Kuang et al. in [11], first builds a $20 \times |X|$-dim profile [8] P_X and then derives a similar $|\Sigma|^k$-dimensional representation from P_X. Such profile representations have been shown to perform well in protein sequence analysis [11, 24]. Constructing the profile for each sequence may not be practical in some application domains, since the size of the profile is dependent on the size of the alphabet set, as well as the difficulty of defining a general sequence similarity search algorithm (e.g., as PSI-BLAST) for non-biological sequence domains. While for bio-sequences $|\Sigma| = 4$ or 20, for music or text classification $|\Sigma|$ can potentially be very large, on the order of tens of thousands of symbols.

The existing string kernel methods essentially amount to the analysis of univariate (i.e. one-dimensional) sequences over finite alphabets Σ with one-dimensional k-mers as basic sequence features. However, original input sequences are often in the form of *sequences of feature vectors*, i.e. each input sequence X is a *sequence of identically sized (R-dim) feature vectors* which could be considered as a $R \times |X|$ feature matrix.

Examples of these multivariate feature sequences include

- *Music data.* A music sequence X in the commonly used MFCC feature representation [16, 20] is a sequence of 13-dimensional MFCC feature vectors, i.e. a multivariate sequence of size $13 \times |X|$.
- *Image data.* An image can be considered as a multivariate sequence of feature vectors extracted from image patches (e.g., as in [23]);
- *Biological data.* Protein sequences can be viewed as *profiles* [11], or as multivariate sequences of R-dim feature vectors describing physical/chemical properties of individual amino acids [30].

While typical string kernel methods essentially use *symbolic Hamming distance-based* matching (as in Eq. 2), recent work in [13] introduced the so-called generalized similarity (*non-Hamming*) kernels that allow to incorporate general similarity metrics $\mathcal{S}(\cdot,\cdot)$ into similarity evaluation and improve performance compared to symbolic Hamming distance-based matching [13]. In particular, most related to the current work, are the distance-preserving symbolic embedding kernels [13] which use similarity hashing [35] to obtain *binary* representations for sequences such that Hamming distance $h(\cdot,\cdot)$ between these binary representations is proportional to the original similarity score $\mathcal{S}(\cdot,\cdot)$ [13]. In contrast, the direct feature quantization (DFQ) method proposed in this work results in more accurate *non-binary* representations that are simpler as they do not require Hamming embedding learning step as in [13], and display higher accuracy (see Experiments, Sec. 4) compared to the binary Hamming-based distance-preserving embedding.

Related methods for the time series classification have also been introduced and include a large variety of methods, e.g., kernels on dynamical systems [33], or alignment-based methods [6, 28] with a quadratic time complexity. We will compare in the experiments with a number of these methods for time series.

In this work, in contrast to kernel methods on univariate string representations, we aim at methods that *directly* exploit multivariate sequence representations to improve accuracy and propose a family of efficient, linear-time *discrete multivariate* similarity kernels (MVDFQ-SK, MVDFQM-SK) using direct feature quantization (DFQ) and manifold kernel embedding (Sec. 3.1, 3.3). We show empirically (Sec. 4) that proposed MVDFQ/MVDFQM kernels and manifold embedding (Sec. 3.3) provide effective improvements in practice over traditional univariate (1D) VQ-based sequence kernels, binary similarity-preserving embedding kernels [13], as well as other state-of-the-art sequence classification methods for a number of challenging classification problems.

3. MULTIVARIATE DIRECT FEATURE QUANTIZATION METHOD

In a typical sequence classification setting, string kernels are restricted to the univariate string data, e.g., word sequences in text analysis, amino acid sequences in the biological sequence analysis, or codeword sequences in the time series analysis [23, 22, 14, 19].

In order to apply these univariate kernel functions to multivariate (R-dimensional) sequences, individual feature vectors at each position in the sequence in the widely used *codebook learning framework* are first encoded using codebook IDs (Figure 1), then standard univariate string kernel methods can be applied on these discrete codeword sequence representations (see e.g., [23, 14]).

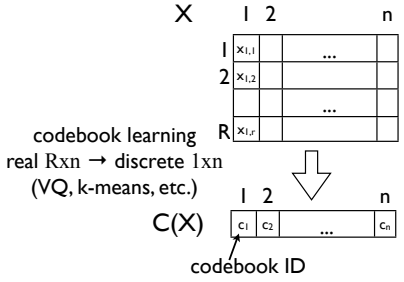

Figure 1: A typical codebook-based representation for a R-dim sequence X: R-dimensional input feature vectors are encoded using corresponding codebook IDs. Univariate string kernels are used to compute similarity between sequences (e.g., [23, 14])

As illustrated in the Fig. 1, the R-dim features vectors from input sequences are first quantized (clustered) to obtain a codebook $\mathcal{C}$, a set of codebook (prototype) vectors, $\mathcal{C} = \{C_1, C_2, ..., C_N\}$, for instance, by applying a Vector Quantization (VQ) algorithm. Then a multivariate input sequence X, a sequence of $n = |X|$ identically sized (real-valued) R-dimensional feature vectors,

$$\mathbf{X} = (x_1, x_2, \ldots, x_n), x_i \in \mathcal{R}^R \; \forall i$$

is encoded as a *univariate (1D) discrete sequence* $c(X)$ of codebook IDs

$$\mathbf{c}(x) = (c_1, c_2, \ldots, c_n), c_i \in \{1 \ldots D\} \; \forall i$$

by mapping each of the vectors x_i to the nearest codeword vector c_i in the codebook $\mathcal{C} = \{C_1, C_2, \ldots, C_D\}$. The resulting codeword sequence $\mathbf{c}(x)$ is essentially a discrete sequence over finite alphabet $\Sigma = \{1, \ldots, D\}$. Univariate (1D) string kernels can then be used for classification with SVM.

In contrast to these commonly employed codebook-based univariate representations, in this work we consider an alternative *multivariate* direct quantization which preserves feature information for each dimension with the manifold embedding representations (MVDFQ/MVDFQM) of the original multivariate (continuous-valued) sequences.

In the following, we first describe the direct feature quantization (DFQ) representations (Sec. 3.1) and contrast them with vector quantization / codebook based representations. We then define a novel family of kernels on these multivariate DFQ representations (Sec. 3.2).

3.1 Direct feature quantization

A multivariate direct feature quantization (MVDFQ) representation of the original continuous-valued multivariate sequence

$$X = (x_1, x_2, \ldots, x_n), x_i \in \mathcal{R}^R$$

is obtained by the direct quantization of each of the R feature dimensions. In this approach, each jth feature $f^j, j = 1 \ldots R$ is quantized by dividing its range (f^j_{min}, f^j_{max}) into the finite number of intervals, B.

In the simplest case, the intervals can be defined, for instance, using a uniform quantization with a pre-specified number of bins B, where the entire feature data range is divided into B equal intervals of length $\delta = (f_{max} - f_{min})/B$ and the index of the quantized feature value $Q(f) = \lfloor (f - f_{min})/\delta \rfloor$ is used to represent the feature value f.

Partitioning of the feature data range could also be obtained by using 1D clustering, e.g., k-means, to adaptively choose dicretization levels and the number of bins for each dimension. Discretization levels also can be chosen using, for example, Gaussian distribution assumption (see, e.g., [21]) as breakpoints under Gaussian curve producing equal-sized areas.

Varying the number of quantization levels B will result in more accurate (larger B) or more coarse (smaller B) representation of the original real-valued data. Here we choose appropriate number of quantization levels B using a small scale cross-validation experiments on the subset of the training data.

Figure 2 shows an example of a DFQ representation for the 3-dimensional time series X (R=3) where the 3-dimensional DFQ representation has been obtained using a uniform binning (B=64) along each of the three data dimensions. As can be seen from the figure, compared to the vector quantization approach, the DFQ retains feature values along each dimension, thus providing a more accurate description of the original real-valued sequence.

We will show in the experiments that using DFQ *multivariate representations* and MVDFQ kernels described below can *significantly (by 25-40%) improve* predictive accuracy compared to traditional 1D (univariate) kernel representations as well as other state-of-the-art approaches (Sec. 4).

3.2 Multivariate Direct Feature Quantization Similarity Kernels

In the following, we first define an efficient multivariate DFQ similarity kernel (MVDFQ-SK) $K(DFQ(X), DFQ(X))$ for the direct feature quantization (DFQ) representation defined in Sec. 3.1. We then present MVDFQ with the man-

Original multivariate time series X

	t=1	t=2	t=3	t=4	t=5
dim 1	0.43	1.43	3.79	2.53	3.29
dim 2	-0.34	0.91	2.97	1.68	2.12
dim 3	-0.41	0.40	2.22	1.15	1.74

Discrete multivariate representation (DFQ(X), B=64 bins)

	t=1	t=2	t=3	t=4	t=5
dim 1	37	38	40	39	39
dim 2	20	21	22	21	21
dim 3	31	32	33	32	33

Univariate representation (VQ(X), codebook size=2048)

	t=1	t=2	t=3	t=4	t=5
dim 1	309	173	484	1148	1252

Figure 2: The proposed discrete multivariate representation (DFQ). The discrete representation DFQ(X) is obtained from the original R=3-dimensional continuous-valued multivariate sequence X by directly discretizing each of R feature dimension. In contrast, the typically used vector quantization approach (codebook) represents multivariate sequence as a one-dimensional discrete sequence of codeword indices.

ifold embedding (MVDFQM) in Sec. 3.3 that as we show experimentally further improve predictive ability of the classifiers on a number of challenging tasks and datasets.

To compute similarity between two multivariate sequences X and Y, we propose a kernel function defined as

$$K_{MVDFQ}(DFQ(X), DFQ(Y)) = \sum_{\alpha_{R\times k}\in DFQ(X)} \sum_{\beta_{R\times k}\in DFQ(Y)} \mathcal{K}(\alpha_{R\times k}, \beta_{R\times k}) \quad (4)$$

where $\alpha_{R\times k}$ and $\beta_{R\times k}$ are $R \times k$ submatrices contained in $DFQ(X)$ and $DFQ(Y)$ and $\mathcal{K}(\alpha_{R\times k}, \beta_{R\times k})$ is a kernel function defined for measuring similarity between two $R\times k$ submatrices. Similarly to k-mer based kernel functions (e.g., Eq. 2), this kernel function essentially computes the similarity between sequences by a cumulative comparison of all pairs of $R \times k$ submatrices contained in $DFQ(X)$ and $DFQ(Y)$ using a submarix kernel function $\mathcal{K}(\cdot,\cdot)$.

One natural definition for the submatrix kernel $\mathcal{K}(\cdot,\cdot)$ is cumulative *row-based comparison*

$$\mathcal{K}(\alpha_{R\times k}, \beta_{R\times k}) = \sum_{r=1}^{R} I_{1\times k}(\alpha^r_{R\times k}, \beta^r_{R\times k}) \quad (5)$$

where $I_{1\times k}(\cdot,\cdot)$ is a similarity/indicator function for matching 1D rows $\alpha^r_{R\times k}$ and $\beta^r_{R\times k}$. The matching function $I_{1\times k}(\cdot,\cdot)$ could be defined as $I_{1\times k}(\alpha,\beta) = 1$ if $d(\alpha,\beta) \le m$, and 0 otherwise (similar to the mismatch kernel).

In the experiments, we use the state-of-the-art spectrum, mismatch [19], and spatial sample (SSSK) [14] kernel functions as our one-dimensional row matching function

$$I_{1\times k}(\alpha^r_{R\times k}, \beta^r_{R\times k})$$

in Eq. 5, which results in corresponding *multivariate* DFQ spectrum, mismatch, and spatial sample kernels (referred as MVDFQ-Spectrum, MVDFQ-Mismatch, and MVDFQ-SSSK, respectively).

Intuitively, according to the kernel definition (Eq. 5), similar $R \times k$ submatrices (i.e. submatrices with many similar rows) will result in a large kernel value $\mathcal{K}(\cdot,\cdot)$.

Using Eq. 5, the multivariate DFQ kernel in Eq. 4 can be written as

$$K_{MVDFQ}(X,Y) = \sum_{r=1}^{R} \sum_{\alpha_{R\times k}\in DFQ(X)} \sum_{\beta_{R\times k}\in DFQ(Y)} I_{1\times k}(\alpha^r_{R\times k}, \beta^r_{R\times k}) \quad (6)$$

which can be efficiently computed by running the corresponding kernel with a 1D k-mer matching function $I_{1\times k}(\cdot,\cdot)$ B times, i.e. for each row $b = 1 \dots R$. The overall complexity of evaluating multivariate kernel $K_{MVDFQ}(DFQ(X), DFQ(Y))$ for two R-dim DFQ sequences $DFQ(X)$ and $DFQ(Y)$ is then $O(R\cdot k\cdot n)$, i.e. is *linear* in the sequence length $n = |X|$ and the number of dimensions R.

3.3 Manifold embedding

While typical string kernel methods assume Euclidean feature space and use Euclidean distance, a *probabilistic manifold assumption* on the geometry of the data space could be more natural and effective (see e.g. [38, 10]). Given d-dim feature representation of a sequence, $\Phi(X) = (\phi_1(X), \dots, \phi_d(X))$, the sequence X can be considered as a point on the multinomial manifold using L1 embedding of $\Phi(X)$:

$$\hat{\Phi}(X) = \left(\frac{\phi_1(X)}{\sum_i \phi_i(X)}, \dots, \frac{\phi_d(X)}{\sum_i \phi_i(X)}\right) \quad (7)$$

where, e.g., in the simple k-mer frequency representation, $\phi_i(X) = f(k_i, X)$, the frequency of k-mer k_i in sequence X.

Then, a natural measure of affinity between the distributions $\hat{\Phi}(X)$ and $\hat{\Phi}(Y)$ on the multinomial manifold is a Bhattacharyya affinity [2], i.e.

$$K_{manifold}(X,Y) = < \sqrt{\hat{\Phi}(X)}, \sqrt{\hat{\Phi}(Y)} > = \sum_i \sqrt{\hat{\phi}_i(X)}\sqrt{\hat{\phi}_i(Y)}$$

Using the equation above and equations for MVDFQ Eq. 6, we obtain the MVDFQ kernel with the *manifold* embedding (MVDFQM):

$$K_{MVDFQM}(X,Y) = \sum_{r=1}^{R} \sum_{\gamma} \sqrt{\phi^r_\gamma(X)}\sqrt{\phi^r_\gamma(Y)} \quad (8)$$

where $\gamma \in \{1, \dots, B\}^K$ is a k-mer over the discretization

alphabet $\Sigma = \{1, \ldots, B\}$ and

$$\phi^r_\gamma(X) = \sum_{\alpha^r_{R\times k} \in DFQ(X)} I(\alpha^r_{R\times k}, \gamma) \quad (9)$$

is the number of occurrences of γ in the r-th dimension/row of X.

In the experiments, we test the manifold embedding with MVDFQ kernel as well as other standard string kernels (we will refer to the MVDFQ with the manifold embedding as MVDFQM, and to the standard VQ kernels with the manifold embedding as VQ-M).

3.4 Advantages of multivariate DFQ

The proposed multivariate DFQ kernel method has the following merits:

- It improves the predictive ability of typical discrete univariate kernel methods with VQ by applying them jointly to multiple discrete sequences obtained from direct discretization of each data dimension of the original real-valued multidimensional sequence.
- Unlike the state-of-the-art approach of quantizing high-dimensional data samples into codewords, it allows for classifier to learn importance of each feature for classification, as the significance of each data dimension for classification can be different.
- It does not rely on clustering or binary similarity-preserving hashing techniques (e.g., as in [13]) as it directly discretizes the feature space using, e.g., uniform binning or adaptive clustering algorithm (k-means).
- It has a low computational cost as it runs in linear time and is scalable to large sequence data sets.
- It can be used with any of the existing univariate sequence kernels (mismatch/spectrum [19], kernels [14], gapped/subsequence kernels [22, 17], etc) to improve performance.

4. EXPERIMENTAL EVALUATION

We study the performance of our methods in terms of the predictive accuracy and the running time on a number of challenging sequence classification problems using standard benchmark datasets for the music genre classification and artist recognition, as well as protein sequence analysis.

4.1 Datasets and experimental setup

We test proposed methods on a number of multi-class sequence classification tasks:

1. 10-class music genre classification[1]. This dataset is a reference music genre recognition dataset introduced in [31, 20]. It contains 1000 30-sec song fragments grouped into 10 genres (blues, rock, classical, etc.), with each genre represented by 100 songs. The task here is to correctly predict the genre of the musical sample.
2. 6-class music genre recognition (ISMIR contest[2]). This is a benchmark music genre recognition task with samples classified into 6 genres.
3. 20-class music artist identification (artist20 dataset[3]). This benchmark dataset contains songs from 120 albums (6 albums per artist) with the task of correctly identifying artist for songs from previously unseen albums.
4. protein remote homology detection (7329 sequences, 54 experiments) [36, 24]. The task here is to correctly infer membership of a given protein in protein superfamilies.

Table 1 provides details of the datasets used in the experiments.

For all music classification tasks input sequences are multivariate sequences of 13-dimensional MFCC feature vectors.

4.2 Baseline methods

We compare the proposed multivariate direct feature quantization (MVDFQ) kernel approach for the multivariate sequence classification to three related baselines:

1. Traditional vector quantization (VQ) approaches with univariate string kernels. This approach has been used in a number of previous studies for the image categorization, text analysis, music classification (see e.g., [23, 15, 14]).
2. A similarity hashing-based kernel approach described recently in [13] (a Euclidean similarity-preserving binary Hamming embedding of original real-valued MFCC feature vectors).
3. A multivariate VQ approach using multiple codebooks of different sizes. In this approach, VQ representations are stacked to obtain essentially a multivariate VQ codebook sequence representation.

We also compare with a number of other state-of-the-art methods specifically developed for the music sequence classification, namely multivariate autoregressive models [26], multilinear models [25], as well as methods with more problem-specific and sophisticated features (aggregate Adaboost [1], classifier fusion with rich spectral and cepstral features [16], non-negative matrix factorization-based approaches [9]).

We test our methods using the state-of-the-art spectrum/mismatch [17] and spatial (SSSK) [14] kernels as our basic univariate kernels, i.e. to implement row matching functions $I_{1\times k}(\cdot,\cdot)$ in Eq. 4).

We use Support Vector Machines (SVMs) classifiers with all kernels.

We also explore the performance impact of varying the number of codewords, discretization bins, a discretization algorithm (uniform, k-means).

4.3 Evaluation measures

For music genre classification experiments a standard 5-fold cross-validation procedure is used as in previous studies [20, 1] to evaluate classification performance. For music artist recognition, we follow a 6-fold leave-one-album-out validation procedure proposed in the previous work [7].

We report average multi-class classification errors as well as F1 scores for all tasks.

[1] http://opihi.cs.uvic.ca/sound/genres

[2] http://ismir2004.ismir.net/genre_contest/index.htm

[3] http://labrosa.ee.columbia.edu/projects/artistid/

Table 1: Benchmark datasets

Dataset	# seqs	# classes	Evaluation
Music genre [31, 20]	1000	10	5-fold cross-validation error
ISMIR2004 contest	1458	6	5-fold cross-validation error
Artist [7]	1413	20	6-fold hold-one-album-per-artist validation error
SCOP protein homology [36]	7329	54	54 binary remote homology detection tasks

Table 2: Music genre recognition (10-class). Comparison of the error rates (%) for the vector quantization (VQ), similarity hashing [13], and the proposed multivariate direct feature quantization (MVDFQ). All methods are compared with (the right three columns) and without (three columns on the left) the proposed manifold embedding (Sec. 3.3). MVDFQ performs better than traditional VQ kernels and sim. hashing kernels [13]

Method	VQ	Sim. hashing [13]	MVDFQ	VQ-M	Sim. hashing-M	MVDFQM
Spectrum	34.5	29.2	28.5	26.9	24.7	23.0
Mismatch	32.6	29.7	24.7	31.2	23.4	20.3
SSSK	31.1	25.9	**22.6**	25.0	23.0	**17.3**

4.4 Parameters and settings

For the vector quantization (VQ) models, we construct codebooks with 2048 codewords from input MFCC vectors. For the multiple-codebook VQ method, we use codebooks with 1024, 2048, an 4096 codewords.

We test our direct feature quantization approach using (1) a uniform quantization of each feature dimension into a fixed number of bins (B=32) and (2) using k-means clustering along each dimension to adaptively select quantization levels and the number of bins per dimension. The number of bins have been found from an initial cross-validation on the subset of training data.

During the testing (classification), for input values outside of the (f_{min}, f_{max}) range, we use special values of 0 and B+1 for values smaller than f_{min} or larger than f_{max}.

For the discrete embedding with similarity hashing [13], we set the number of bits $E = 32$ which has been found to perform well in [13].

The length of k-mers used in the MVDFQ-spectrum/mismatch kernels has been set to k=6 and and the number of mismatches is set as m=1 for the mismatch kernels (these values have been selected using cross-validation on the subset of the training data).

For the univariate kernels on codeword (VQ) settings, the best settings of the k-mer length and the number of mismatches are similar, $k = 5 - 6$ and m=1,2.

For the spatial sample kernels (SSSK) [14] we use the spatial kernel with three (t=3) k=1-mer features and the maximum distance parameter d=5.

All experiments are performed on a single $2.8GHz$ CPU. The datasets used in our experiments and the supplementary data/code are available at
`http://pkuksa.org/~pkuksa/mvdfq.html`.

4.5 Music genre recognition

We first compare the proposed multivariate DFQ kernel approaches (MVDFQ) (Sec. 3.2, 3.3) with the vector quantization-based univariate kernels, and the recently proposed binary Hamming similarity hashing kernels [13].

All of the methods are compared with and without the proposed manifold embedding (Sec. 3.3).

As shown in Table 2, on a widely used benchmark dataset for music genre recognition [20, 31] (10 genres, each with 100 sequences), proposed *multivariate* DFQ kernels improve over traditional univariate VQ kernels, as well as recently proposed binary Hamming similarity hashing kernels [13]. Using direct feature quantization (DFQ) kernels effectively improves accuracy compared to the VQ, and the similarity hashing approach for all basic kernels (spectrum, mismatch, and sparse spatial sample kernels (SSSK)). For instance, MVDFQ-SSSK achieves a significantly lower error rate of 22.6% compared to 31.1% using VQ or 25.9% using similarity hashing (27% and 13% improvements, respectively).

Using the k-means clustering for the discretization with DFQ results in the performance similar to the uniform quantization (e.g., MVDFQ-Spectrum and MVDFQ-SSSK with k-means achieve slightly lower errors of 22.8% and 16.9% compared to 23.0% and 17.2% with the uniform quantization).

We also note that using the manifold embedding further reduces error for all of the methods, including the VQ-based and similarity hashing kernels. These consistent improvements in accuracy across all of the methods, could be attributed to the manifold embedding effectively exploiting intrinsic geometric structure of music data. Overall, MVDFQ with manifold embedding (MVDFQM-SSSK) achieves the best error rate of 17% compared to the the best 23% error rate of similarity hashing kernel or 25.0% using VQ. (24% and 30% relative improvements in the error rates, respectively).

We also compare with previous best results on this music genre recognition dataset and baselines in Table 3, including multivariate autoregressive models [26], wavelet-based DWCH [20] method, aggregate AdaBoost [1], approaches specifically developed for the music classification that also use many other features in addition to MFCC. As can be seen from the table, using multivariate DFQ kernels (MVDFQ-SK) compares well with the state-of-the-art results (e.g., AdaBoost method [1] with much richer feature set).

The proposed MVDFQM method is also more effective than using the multiple codebook VQ method. We also note that expanding feature set by adding a set of 64 FFT features to the 13 MFCC feature set (i.e. $77 \times |X|$ multivariate representation), could further increase accuracy to 86.4% compared to that of 82.7% with MFCC features alone (Table 3).

Table 3: Music genre recognition (10-class). Comparison with previous results and baselines.

method	Error, %	F1
Baseline 1: MFCC (Mean+Variance)	48.30	51.55
Baseline 2: NMF [9]	26.0	-
Baseline 3: Multiple codebook VQ	29.6	70.51
Baseline 3: (non-MFCC): DWCH [20]	21.5	-
Baseline 4: MAR (multivariate autoregressive model) [26]	21.7	-
Baseline 5: AdaBoost (MFCC,FFT,LPC,etc) [1]	17.5	-
MVDFQ-SSSK (MFCC)	26.3	73.88
MVDFQM-SSSK (MFCC)	17.3	82.85
MVDFQM-SSSK (MFCC+FFT64)	**13.6**	**86.53**

We also note the utility of the *multivariate* direct feature quantization representation and kernel (MVDFQ) as opposed to the *univariate* representations: the univariate kernel on the one-dimensional (1D) sequence obtained from the DFQ multivariate sequence by encoding each R=13-dim feature vector as one codeword $c = \sum_i^R Q(f_i)B^{i-1}$, i.e. using alphabet size $|\Sigma| = B^R$, gives a higher error rate of 22% compared to that of 17.3% when using the *multivariate* DFQ kernel with the manifold embedding (MVDFQM).

We observe similar overall improvements for multivariate MVDFQ-SK kernels on another benchmark dataset (ISMIR2004 genre contest), Table 4. For instance, using the MVDFQ-SK string kernel with uniform direct feature quantization (MFCC only) reduces the error rate to 16.7% compared to that of 19.6% when using VQ with the univariate kernel.

As also can be seen from the Table 5, obtained error rates (16-17%) for ISMIR genre recognition compare well with a number of previous best results, including the recent non-negative matrix factorization method [9](16.5% error), or complex auditory model and cortical representation of [25] (19% error).

4.6 Artist recognition

We also illustrate the utility of our multivariate MVDFQ-SK kernels and representations on the multi-class artist identification problem using the standard *artist20* dataset with 20 artists, 6 albums each (1413 tracks total). Table 6 lists results for the 6-fold album-wise cross-validation with one album per artist held out for testing. Using multivariate MVDFQ-SK kernels with the direct uniform quantization of MFCC features yields a much lower 25.7% error compared to the best error rate for 42.9% for univariate kernels (a 40% relative improvement in error).

4.7 Protein remote homology detection

In Table 7, we compare our proposed *multivariate* DFQ string kernel method (using 20-dim BLOSUM rows as feature vectors for individual amino acids, i.e. $20 \times |X|$ multivariate sequences) with a number of state-of-the-art kernel methods for the remote homology detection including spectrum/mismatch kernels [19, 18], spatial sample kernels [14], similarity hashing kernels [13], as well as recently proposed spectrum-RBF and mismatch-RBF methods [30] which also incorporate *physico-chemical descriptors*.

As can be seen from results in Table 7, multivariate DFQ string kernel (MVDFQ) provides effective improvements over other methods. For instance, using MVDFQ spectrum and mismatch kernels with BLOSUM substitution profiles significantly improves average ROC50 scores from 27.91 and 41.92 to 43.29 and 49.17, respectively (relative improvements of 50% and 17%), compared to traditional univariate spectrum/mismatch approaches.

4.8 Running time

In Table 8, we compare the classification performance and the running time of our method to the recent binary Hamming embedding (similarity hashing) [13], and vector quantization-based univariate kernels. We vary the dimensionality of the embedding space E, the codebook size, and the number of discretization bins B, respectively. We note that for mismatch-(k,m) kernel computation we use the linear time sufficient-statistic based algorithm from [12]) and for all other methods we use their existing state-of-the-art implementations.

As can be seen from Table 8, multivariate kernels with the direct feature quantization (MVDFQ) display a better performance compared to the similarity hashing [13] and traditional univariate kernels.

5. CONCLUSIONS

We presented novel discrete multivariate direct feature quantization kernel methods (MVDFQ-SK and MVDFQ-SK with the manifold embedding) for data in the form of sequences of feature vectors (as in music MFCC sequences, biological sequence profiles, or image sequences). The proposed approach directly exploits original multivariate feature sequences to improve sequence classification as opposed to using univariate codeword sequences. On three music classification tasks as well as protein sequence classification this shows significant 25-40% improvements compared to the traditional codebook learning and state-of-the-art sequence classification methods.

6. REFERENCES

[1] J. Bergstra, N. Casagrande, D. Erhan, D. Eck, and B. Kégl. Aggregate features and adaboost for music classification. *Mach. Learn.*, 65:473–484, December 2006.

[2] A. Bhattacharyya. On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society*, 35:99–109, 1943.

[3] W. M. Campbell. Generalized linear discriminant sequence kernels for speaker recognition. In *Acoustics, Speech, and Signal Processing (ICASSP), 2002 IEEE International Conference on*, volume 1, pages I–161 –I–164, may 2002.

Table 4: Comparison of the vector quantization (VQ), the similarity hashing [13] approach, and the proposed multivariate DFQ method (MVDFQ) on the music genre recognition ISMIR contest, error rates (%)

method	Error, %	**method** (with Manifold)	Error, %
VQ	24.15	VQ-M	19.62
Sim. hashing [13]	22.63	Sim. hashing-M	17.83
MVDFQ	**19.89**	MVDFQM	**16.74**

Table 5: Classification performance on the ISMIR contest data

method	Error, %	F1
Baseline 1: MFCC (Mean+Variance)	34.85	58.15
Baseline 2: Multilinear (Cortical) [25]	19.05	-
Baseline 3: pLSA [37]	19.9	-
Baseline 4 (best): NMF [9]	16.5	-
MVDFQM-Spectrum (MFCC)	**16.74**	**80.57**
MVDFQM-Spectrum (MFFC+FFT64)	**16.19**	**80.79**

Table 6: Music artist identification (20-class, 13-dim MFCC features only)

method	Error, %	F1
Baseline 1: MFCC (Mean+Variance)	59.54	40.91
Baseline 2: MFCC+FFT64 (Mean+Variance)	54.17	47.09
Baseline 3: MFCC (GMM) [7]	44.0	-
VQ	42.97	57.26
Sim. hashing [13]	34.62	66.22
MVDFQM-Spectrum	**25.67**	**74.79**

Table 7: Classification performance (mean ROC50) on protein remote homology detection (54 experiments). Multivariate DFQ string kernels perform better than traditional (univariate) kernels and similarity-preserving hashing kernels [13] (p-values indicate statistical significance of the observed differences between univariate and corresponding multivariate DFQ kernels)

	Mean ROC50			
Method	Univariate	Sim. hashing [13]	Multivariate DFQ	p-value
Spectrum [18]	27.91	35.14	**43.29**	2.5e-6
Mismatch-(k=5,m=1) [19]	41.92	46.35	**49.17**	6.5e-3
Spatial sample (SSSK) [14]	50.12	52.0	**55.54**	2.7e-5
Spectrum-RBF [30]	42.1	-	-	-
Mismatch-RBF [30]	43.6	-	-	-

Table 8: Music genre classification performance and running time for the kernels using the vector quantization, similarity hashing [13], and the proposed direct feature quantization (DFQ) approach as a function of the codebook size, the embedding size, and the number of discretization bins, respectively

Embedding size	Error, %	Running time (s), 1000 × 1000 kernel matrix computation
Mismatch(5,2) VQ $\|\Sigma\|$=2048	32.6	28
Mismatch(5,2) VQ $\|\Sigma\|$=1024	32.5	26
Sim. hashing E=16,k=5	28.3	24
Sim. hashing E=32,k=5	24.7	61
Sim. hashing E=64,k=5	28.2	202
MVDFQ B=16	23.9	9.4
MVDFQ B=32	23.0	10.4
MVDFQ B=64	22.8	11.6

[4] J. Cheng and P. Baldi. A machine learning information retrieval approach to protein fold recognition. *Bioinformatics*, 22(12):1456–1463, June 2006.

[5] C. Cortes, P. Haffner, and M. Mohri. Rational kernels: Theory and algorithms. *Journal of Machine Learning Research*, 5:1035–1062, 2004.

[6] M. Cuturi, J.-P. Vert, Ø.. Birkenes, and T. Matsui. A kernel for time series based on global alignments. In *Proceedings of the Intern. Conference on Acoustics, Speech and Signal Processing*, volume II, pages 413 – 416, 2007.

[7] D. P. W. Ellis. Classifying music audio with timbral

and chroma features. In S. Dixon, D. Bainbridge, and R. Typke, editors, *ISMIR*, pages 339–340. Austrian Computer Society, 2007.

[8] M. Gribskov, A. McLachlan, and D. Eisenberg. Profile analysis: detection of distantly related proteins. *Proceedings of the National Academy of Sciences*, 84:4355–4358, 1987.

[9] A. Holzapfel and Y. Stylianou. Musical genre classification using nonnegative matrix factorization-based features. *Audio, Speech, and Language Processing, IEEE Transactions on*, 16(2):424 –434, feb. 2008.

[10] T. Jebara, R. Kondor, and A. Howard. Probability product kernels. *J. Mach. Learn. Res.*, 5:819–844, December 2004.

[11] R. Kuang, E. Ie, K. Wang, K. Wang, M. Siddiqi, Y. Freund, and C. S. Leslie. Profile-based string kernels for remote homology detection and motif extraction. In *CSB*, pages 152–160, 2004.

[12] P. Kuksa, P.-H. Huang, and V. Pavlovic. Scalable algorithms for string kernels with inexact matching. In *NIPS*, 2008.

[13] P. P. Kuksa, I. Khan, and V. Pavlovic. Generalized similarity kernels for efficient sequence classification. In *SDM*, 2012.

[14] P. P. Kuksa and V. Pavlovic. Spatial representation for efficient sequence classification. In *ICPR*, 2010.

[15] P. P. Kuksa and Y. Qi. Semi-supervised bio-named entity recognition with word-codebook learning. In *SDM*, 2010.

[16] C.-H. Lee, J.-L. Shih, K.-M. Yu, and H.-S. Lin. Automatic music genre classification based on modulation spectral analysis of spectral and cepstral features. *Multimedia, IEEE Transactions on*, 11(4):670 –682, june 2009.

[17] C. Leslie and R. Kuang. Fast string kernels using inexact matching for protein sequences. *J. Mach. Learn. Res.*, 5:1435–1455, 2004.

[18] C. S. Leslie, E. Eskin, and W. S. Noble. The spectrum kernel: A string kernel for SVM protein classification. In *Pacific Symposium on Biocomputing*, pages 566–575, 2002.

[19] C. S. Leslie, E. Eskin, J. Weston, and W. S. Noble. Mismatch string kernels for SVM protein classification. In *NIPS*, pages 1417–1424, 2002.

[20] T. Li, M. Ogihara, and Q. Li. A comparative study on content-based music genre classification. In *SIGIR '03*, pages 282–289, New York, NY, USA, 2003. ACM.

[21] J. Lin, E. Keogh, S. Lonardi, and B. Chiu. A symbolic representation of time series, with implications for streaming algorithms. In *Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, DMKD '03, pages 2–11, New York, NY, USA, 2003. ACM.

[22] H. Lodhi, C. Saunders, J. Shawe-Taylor, N. Cristianini, and C. Watkins. Text classification using string kernels. *J. Mach. Learn. Res.*, 2:419–444, 2002.

[23] Z. Lu and H. Ip. Image categorization with spatial mismatch kernels. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 0:397–404, 2009.

[24] I. Melvin, E. Ie, J. Weston, W. S. Noble, and C. Leslie. Multi-class protein classification using adaptive codes. *J. Mach. Learn. Res.*, 8:1557–1581, 2007.

[25] I. Panagakis, E. Benetos, and C. Kotropoulos. Music genre classification: A multilinear approach. In *ISMIR'08*, pages 583–588, 2008.

[26] H. Rump, S. Miyabe, E. Tsunoo, N. Ono, and S. Sagayama. Autoregressive mfcc models for genre classification improved by harmonic-percussion separation. In J. S. Downie and R. C. Veltkamp, editors, *Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR 2010, Utrecht, Netherlands, August 9-13, 2010*, pages 87–92. International Society for Music Information Retrieval, 2010.

[27] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, New York, NY, USA, 2004.

[28] H. Shimodaira, K.-I. Noma, M. Nakai, and S. Sagayama. Dynamic time-alignment kernel in support vector machine. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, Cambridge, MA, 2002. MIT Press.

[29] S. Sonnenburg, G. Rätsch, and B. Schölkopf. Large scale genomic sequence SVM classifiers. In *ICML '05*, pages 848–855, New York, NY, USA, 2005.

[30] N. Toussaint, C. Widmer, O. Kohlbacher, and G. Ratsch. Exploiting physico-chemical properties in string kernels. *BMC Bioinformatics*, 11(Suppl 8):S7, 2010.

[31] G. Tzanetakis and P. Cook. Musical genre classification of audio signals. *Speech and Audio Processing, IEEE Transactions on*, 10(5):293 – 302, jul 2002.

[32] V. N. Vapnik. *Statistical Learning Theory*. Wiley-Interscience, September 1998.

[33] S. V. Vishwanathan, A. J. Smola, and R. Vidal. Binet-cauchy kernels on dynamical systems and its application to the analysis of dynamic scenes. *Int. J. Comput. Vision*, 73(1):95–119, June 2007.

[34] S. V. N. Vishwanathan and A. Smola. Fast kernels for string and tree matching. In *NIPS*, 2002.

[35] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 1753–1760. 2009.

[36] J. Weston, C. Leslie, E. Ie, D. Zhou, A. Elisseeff, and W. S. Noble. Semi-supervised protein classification using cluster kernels. *Bioinformatics*, 21(15):3241–3247, 2005.

[37] Z. Zeng, S. Zhang, H. Li, W. Liang, and H. Zheng. A novel approach to musical genre classification using probabilistic latent semantic analysis model. In *Multimedia and Expo, 2009. ICME 2009. IEEE International Conference on*, pages 486 –489, 28 2009-july 3 2009.

[38] D. Zhang, X. Chen, and W. S. Lee. Text classification with kernels on the multinomial manifold. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '05, pages 266–273, New York, NY, USA, 2005. ACM.